

KNOWLEDGE FROM FALSEHOOD

Ted A. Warfield
University of Notre Dame

I am interested in a variety of situations in which one apparently knows something despite the existence of something non-ideal in the epistemic pedigree of the belief in question (examples: faulty reasoning, false premise, problematic testimony). In this paper I begin a discussion of the ‘knowledge from falsehood’ part of this project.

Consider the following claim:

- (1) Inferential knowledge of a conclusion requires known relevant premises.¹

One consequence of (1) is:

- (1a) Inferential knowledge of a conclusion requires true relevant premises.

I will focus my critical attention on (1a).² I will now offer some charitable interpretation of parts of (1a), interpretations that help us dodge obvious counterexamples. I’ll then briefly explain some ways of motivating (1a) before beginning my critical engagement with it.

Clarifications of ‘inferential knowledge’ and ‘relevant’ will enable us to sidestep what might seem to be obvious counterexamples to (1a). I make these clarifications in reverse order. First, the restriction to ‘relevant’ premises avoids simple counterexamples involving overdetermination. I reason (having seen each enter the room a minute ago)—‘Jerry Fodor is in the room, Steve Stich is in the room, Colin McGinn is in the room, Brian McLaughlin is in the room; therefore, at least one Rutgers philosopher of mind is in the room.’ Assume that my premises are all true except one: McGinn has, after entering, stepped out to take a phone call. I know my conclusion despite the false premise about McGinn. Here’s a first pass at an explanation: the McGinn premise is not ‘relevant’ in this case to the particular conclusion. I will address further issues about the proper unpacking of ‘relevant’ when they arise in the discussion of my counterexamples to (1a).

Second, in evaluating (1a) let us charitably understand ‘inferential knowledge’ so that our attention is restricted to cases in which every path to the conclusion in question is an inferential path. Otherwise we would likely find counterexamples in cases where one both sees that P, and infers (through a falsehood) that P, but where one knows because one sees and one’s seeing epistemically dominates the inference. Having been told falsely but authoritatively that Barry Loewer never leaves this room, I reason (through the suppressed conditional ‘If Loewer never leaves this room then he’s here today’) to the conclusion that Loewer is here today. As I complete this reasoning, I see Loewer standing in front of me. I know that Loewer is here despite reasoning through a falsehood to my conclusion.³

In the interest of further simplification, let’s focus on cases in which one has exactly one ‘inferential argument’ for one’s conclusion and one’s inferential argument consists of a single material premise and a suppressed conditional linking the premise to the conclusion via simple modus ponens. This simplification does not require taking a stand on the question of whether coherentist or other ‘holistic’ features play a justificatory role nor does the simplification involve taking a stand on the epistemic role of ‘the background’ or anything else. In simplifying in this way, we simply attend to inferences where the surface form of reasoning involved is quite perspicuous. Simplify the discussion still further by focusing on the *formation* of the belief via this modus ponens inference (and the epistemic status of the belief at the time of formation) and leave belief ‘sustenance’ issues aside.⁴

Turning from clarification to motivation, I now briefly present some ways of motivating (1a). Here are three quick points followed by a fourth and more fully explored point.⁵ First, (1a) partly reflects widespread preference for sound arguments over merely valid arguments. Second, inferential knowledge requires that the conjunction of one’s premises constitutes a ‘good reason’ for accepting one’s conclusion and falsity suffices for the badness of a reason. Third, with a false premise in one’s inferential argument, it appears that luck of some sort is needed to get one to a truth at one’s conclusion.⁶ Compare the luck apparently involved in reaching a true conclusion via valid inference from a *false* premise with the *guarantee* of truth in the conclusion when one’s inferential argument to that conclusion is sound. Turn now to a fourth point that seemingly motivates (1a): a natural line of thought has it that Gettier issues suggest that (1a) is correct.

Begin with some Gettier lessons. I see what looks like a dog out in the yard. I reason (on this occasion) as follows: there’s a dog in the yard, so there’s at least one animal in the yard. As it happens, my premise is false (there’s a realistic looking toy in the yard and that is what I see in the dog-less yard). But my conclusion is true because concealed behind some brush is a squirrel. I don’t know my conclusion despite its truth and the valid argument leading to it. The falsity of my premise is, perhaps, what has gone wrong in the case. In studying the Gettier problem, we learn fairly early on that ‘reasoning through a falsehood’ is not an essential element in constructing a Gettier case. One typically

learns this from a ‘stopped clock case’. At what I correctly judge to be 5 minutes’ walking distance from my classroom, I see a clock face reading 1:00. Relying wholly on that reading of the clock for my estimate of the time, I reason: it’s 1:00; so, I’m not late for my 1:10 class. It is in fact 1:00 but the clock I am reading stopped working days ago. I don’t know I’m not late for my 1:10 class. Reasoning through a relevant falsehood is not necessary for generating Gettier cases. Reflection on the dog/yard/squirrel case, however, leaves many thinking that reasoning through a relevant falsehood is sufficient for generating a Gettier case. Because Gettiered conclusions aren’t knowledge, (1a) would seem to follow.

We thus have non-trivial motivation for (1a). Any rejection of (1a) should surely include some commentary explaining why and how these four motivations for (1a) are safely rejected. I will return to this task after providing and discussing some counterexamples to (1a), and rejecting anticipated resistance to the counterexamples.

I turn now to some examples. (1a) cannot be sustained. Begin with this counterexample.

Example #1. My doctor has ordered that I get at least 8 hours sleep per night. I knowingly go to bed at 11pm. I wake up and see the clock reading “2:30am”. I reason: (premise) I’ve been asleep 3½ hours, so (conclusion) I haven’t slept the mandated 8 hours. My premise is false. I’ve forgotten that it’s the night of the (Fall) time change and I have a clock that automatically resets in the appropriate way at the appropriate time and it has already done this. So I’ve been asleep 4½ hours (we’re at “2:30am” for the 2nd time). Despite all of this complexity (and, indeed, partly because of it) I still know my conclusion.⁷

Additional counterexamples to (1a) are easy to find. Not all of the examples I’ll present are equally compelling and perhaps some are similar enough to others that they add nothing of substance to the investigation. I provide a range of examples so that non-essential features of one example don’t inappropriately dominate the analysis and so that distinct types of counterexamples (if distinct types exist) can perhaps be identified. Here’s another example.

Example #2. With hopes of getting him to attend a party in Providence on Saturday night, Jaegwon Kim asks Christopher Hill what he’s doing on Saturday. Hill replies ‘I’m flying to Fayetteville on Saturday night’ and the conversation ends. Kim, recalling that Hill taught for many years in Fayetteville, Arkansas, reasons as follows: ‘Hill will be in Arkansas on Saturday night; so, he won’t be at my party Saturday night’. Kim knows his conclusion, but his premise is false: Hill is flying to Fayetteville, *North Carolina*.⁸

And here, in abbreviated form, are three more to round out the example set:

Example #3. Counting with some care the number of people present at my talk, I reason: ‘There are 53 people at my talk; therefore my 100 handout copies are

sufficient'. My premise is false. There are 52 people in attendance—I double counted one person who changed seats during the count. And yet I know my conclusion.

Example #4. CNN breaks in with a live report. The headline is 'The President is speaking now to supporters in Utah'. I reason: 'The President is in Utah; therefore he is not attending today's NATO talks in Brussels'. I know my conclusion but my premise is false: the President is in Nevada—he is speaking at a 'border rally' at the border of those two states and the speaking platform on which he is standing is in Nevada. The crowd listening to the speech is in Utah.

Example #5. I have a 7pm meeting and extreme confidence in the accuracy of my fancy watch. Having lost track of the time and wanting to arrive on time for the meeting, I look carefully at my watch. I reason: 'It is exactly 2:58pm; therefore I am not late for my 7pm meeting'. Again I know my conclusion, but as it happens it's exactly 2:56pm, not 2:58pm.

With five apparent counterexamples to (1a) before us, I should clarify the overall dialectical situation and then look at possible responses to the counterexamples from those who don't want to concede the falsity of (1a). The examples are at least apparent counterexamples to (1a) because it seems that:

(i) the examples involve inferential knowledge of a conclusion

and

(ii) the examples involve a false relevant premise.

In defense of (i) there is nothing to do except appeal to clear and widely shared intuitions about the cases. In defense of (ii), the falsity of the premise in each argument is properly stipulated in each example and the 'relevance' of that false premise is suggested by the fact that the premise is the sole material premise in the inferential episode leading to the conclusion. If the sole substantive premise isn't relevant then someone has some explaining to do about the notion of 'relevance' involved in (1a).

In light of the examples and these observations, I prefer the following straightforward evaluation of the situation. We have counterexamples to (1a) and so (1a) is false. The cases are cases of '*knowledge from falsehood*' and so we must integrate this fact, carefully no doubt, into our overall epistemological thinking. As a starting point for this, note that in the examples the truth of the conclusion and the path to it seem quite stable despite the relevant falsehoods in the examples. Metaphorically, we do not seem to be kicked off the path to inferential knowledge of our conclusion in these examples by the falsehood. This seems to be because the falsehood, to leave things unanalyzed for now, seems to do the job (of taking us to knowledge) quite well.

If everyone would agree with this initial reaction to the counterexamples, I would move forward immediately to unpack my metaphorical diagnostic language and to revisit the motivations for believing (1a) to explain how they led us astray. Strangely enough, however, not everyone believes that we should reject (1a). Many would like to resist this conclusion. Any resistance to a counterexample to (1a) must make one of two claims. Either the resistance must claim that the agent in the example does not know the conclusion or the resistance must claim that the falsehood involved in the example is not ‘relevant’ in the sense involved in (1a). Though the examples are not of precisely equal plausibility, I think it is safe to exclude the first approach (denying knowledge) as a *general* response to the range of examples.⁹ The resistance must therefore make at least partial use of the other available response: denying that the involved falsehood is a ‘relevant’ falsehood.

Those favoring this approach should be heartened by the fact that I have not provided an account of (or even a sufficient condition for) the ‘relevance’ of a premise in an inferential argument. In charitably rejecting simple overdetermination counterexamples to (1a), I rejected the simple idea that a sufficient condition for the relevance of a premise is that the premise actually is used in the inferential argument being examined.¹⁰ This gestures at a familiar distinction between the causal and epistemic base of a belief and plausibly suggests that ‘relevance’ is a matter of epistemology. And a moment ago I suggested, but did not *argue*, that in the one premise arguments we are considering, the fact that the false premise is the sole premise seems to suffice for its relevance. The resistance should seize upon this inexactness and attempt to understand the examples in such a way that though knowledge of the conclusion is indeed present, the falsehood involved is not relevant. If successful, the resistance could fairly describe the examples as ‘knowledge *despite* falsehood’ (knowledge despite an involved, though not epistemically relevant, falsehood) rather than ‘knowledge *from* falsehood’ as I would have it.

There are several ways one might attempt to work out a detailed resistance strategy.¹¹ But all possible resistance strategies must make the following claims about each counterexample in which knowledge of the conclusion is conceded:

- (i) though there is a falsehood *involved* in the inferential argument, the falsehood is not *relevant*.
- (ii) a truth (not the falsehood) in the example plays the central epistemizing role.

Point (ii) is the substantive point that downplays the role of the falsehoods in the examples. The falsehoods are involved in the examples, but perhaps they are not epistemically relevant. Let’s see what the resistance might have in mind by returning to the examples. I begin with two clearly inadequate resistance strategies that help point the way to more serious resistance.

Most who have reflected on the five examples displayed earlier have noticed at least the following. In each example there seems to be at least one true

proposition ‘somewhere in the neighborhood’ that seems well suited to playing an epistemizing role with respect to the conclusion. The propositions seem well suited because, at a minimum, it looks like they are strongly supported by the agent’s evidence. Though individuals might identify different propositions in some cases, let’s take note of some candidate resistance propositions (noting possible alternatives in brackets):

Example #1 - I’ve been asleep less than 5 hours. [alternatively ‘about 3 ½ hours’]

Example #2 - Hill will be out of town on Saturday. [‘not in Providence’]

Example #3 - There are approximately 53 people at my talk. [‘fewer than 55’]

Example #4 - The President is in the Western USA. [‘is in or near Utah’]

Example #5 - It’s about 3:00 pm. [‘several hours before my meeting’]

The simplest resistance strategy would, having identified these true propositions in each of the proposed counterexamples, boldly suggest the following response to the counterexamples. The involved falsehoods in the examples are not relevant and so we have no counterexamples to (1a). Epistemization in each example is a product of the above identified propositions. These propositions are all true and the agent in each example seems to have good evidence for the proposition. I know that I haven’t had the doctor-mandated eight hours of sleep, says this approach, not because of my justified belief that I’ve slept 3 ½ hours (that’s false) but because of the nearby true proposition ‘I’ve slept fewer than 5 hours’ that is also justified for me.

This simplest resistance strategy moves too quickly. The strategy suggests that sufficient for knowledge in each case is there being a true proposition intuitively ‘in the neighborhood’ that serves to epistemize each example’s conclusion. This simple approach fails. Notice that though we have identified true propositions intuitively ‘in the neighborhood’ of the examples for which the agents appear to have strong evidence, we have not even said that the agents believe the propositions. One wonders how the propositions identified are supposed to epistemize the conclusions if they are not even believed by the agents in the examples. Given the evidence that the agents seem to have for the identified proposition, it seems plausible that the agents would be disposed to believe the relevant proposition. But mere dispositions to believe cannot play an epistemizing role in an inferential argument: allowing them to do so grossly over-ascribes inferential knowledge.

This first resistance strategy suggests a second. Acknowledging that mere dispositions to believe can’t play an epistemizing role, the resistance might claim that the agent in each example *at least* dispositionally believes the identified proposition. Dispositional beliefs (or stronger) surely can play an epistemizing role. The natural suggestion here is that one knows despite the involvement of a falsehood, if one has a justified dispositional belief suitable for epistemizing one’s conclusion.

This second resistance strategy also clearly fails. Grant for purposes of criticism the controversial assumption that in all cases of the same type as the

proposed counterexamples the agent at least dispositionally believes the identified proposition.¹² Proponents of this strategy encounter severe difficulty with standard ‘basing relation’ cases. The detective believes that Jones is the murderer. He accepts the contents of a forensic report indicating that DNA evidence strongly points to Jones. But the detective believes that Jones is the murderer not on the basis of the report but rather solely because he infers this from his delusional belief that Jones confessed to the killings. The standard and proper verdict in this case is that the detective does not know that Jones is the murderer. The second resistance strategy fails here, delivering the verdict that detective has knowledge. Despite the involved false belief about the confession, there is a true and evidentially supported dispositional belief in the neighborhood (about the contents of the forensic report) suitable for epistemizing the belief that Jones is the murderer. The second resistance strategy therefore says that the detective knows that Jones is the murderer in a case in which uncontroversially he does not know this. This fact eliminates the second resistance strategy from contention.

Once again, however, failure suggests another path. A third resistance strategy can be identified in an attempt to keep what is promising in the second strategy while distinguishing the cases in which we clearly want to ascribe knowledge (the counterexample cases at least) from the cases in which we clearly do not want this to happen (the ‘no knowledge’ standard basing relation cases). Notice that almost all of the truths pointed to in the alleged counterexamples are entailed by the key proposition falsely believed in the example. Notice further that for each example, there is at least one such entailed truth. The resistance might claim that it is the entailment of the justified and (at least) dispositionally believed truth by the involved falsehood that marks these cases as cases of knowledge despite the involvement of a falsehood in the inference.

The suggestion is that one has knowledge despite the presence of an involved falsehood if there is a justified and (at least) dispositionally believed truth entailed by the falsehood that serves as the premise in one’s inferential argument. The truth is the epistemic basis though the falsehood is the causal basis. The truth is an appropriate epistemic basis (avoiding standard basing condition case worries) because of the entailment of the truth by the falsehood. Notice that the detective’s false belief about a confession from Jones does not entail the proposition concerning the content of the forensic report. This shows how this approach would improve upon the second approach’s mishandling of the detective case.

A variant on this proposal can be identified as a fourth possible resistance strategy. The resistance might not be best served by requiring that the dispositional belief being pressed into service as the purported epistemizer be *entailed* by the falsehood. Perhaps the resistance could make do with a less restrictive condition that manages to block the verdict that the resistance gets the detective case wrong. The resistance might suggest the following broader account. One has knowledge despite the presence of an involved falsehood if there is a justified

and (at least) dispositionally believed truth evidentially supported by the evidence for the involved falsehood. Notice that just as with the third resistance strategy, the detective case is not a problem because the evidence for the confession belief does not epistemically support the belief concerning the forensic report's contents.

Here are two apparent advantages this fourth strategy may have over the third. This fourth strategy seems to avoid what resistance fans should think is an unsettling picture of the epistemology of the counterexample cases that emerges from the third resistance strategy. According to the third strategy, knowledge is reached in the conclusion but the falsehood seems to be playing an important epistemizing role: there is good evidence for the falsehood, which then entails the truth, which then epistemizes the conclusion. This picture seems to give the falsehood an important and apparently epistemic role to play, one that should offend the intuitions of fans of the resistance. The fourth strategy avoids this appearance. A second apparent advantage of the fourth resistance proposal is its breadth. This account seemingly covers a wider range of cases than the account requiring entailment: entailment is presumably only one kind of epistemic support relation between evidence and conclusion.

I evaluate (and reject) these more serious third and fourth resistance proposals together. Grant, once again, the assumption that in every counterexample case we will find dispositional beliefs entailed by the involved falsehood or epistemically supported by it. Both proposals fail catastrophically in Gettier cases. I seem to see a dog in the yard. On this occasion I form the belief that there is a dog in the yard and then reason as before to the conclusion that there is at least one animal in the yard. My belief is false (there is no dog, only the toy) and my conclusion though true, because of the squirrel behind the brush, is not known. The third and fourth resistance strategies get this clear 'no knowledge' case wrong. They both imply that I know that there's at least one animal in the yard. After all, there is a justified and dispositionally believed truth that is *both* evidentially supported by and entailed by my false belief that there is a dog in the yard: the truth is 'there is a dog or squirrel in the yard'. The third and fourth resistance strategies both rule that this belief epistemizes my conclusion. I therefore am judged to have knowledge in this case which goes against the standard and correct Gettier verdict. The sophisticated third and fourth resistance strategies fail and with them fails the resistance.

I see no way of maintaining that the knowledge present in the counterexample cases traces entirely to some truth. I think we need to take the cases at face value: they are cases of knowledge *from* falsehood, not cases of knowledge *despite* falsehood. We sometimes gain knowledge via inference from a false premise. And not merely in familiar enough cases where a falsehood is not essentially involved (in a non-epistemizing role). I think we need to accept that relevant falsehoods sometimes play a central epistemizing role in inference. As noted earlier, one taking my position owes responses to the motivations that seem to lead so many to embrace (1a). As we'll see momentarily, these responses

raise complex and important questions about central epistemological issues. I can't deal adequately with all of these issues on this occasion. Here I'll restrict my focus to showing what issues need to be addressed by defending my position and by beginning to sketch my preferred strategy for dealing with these issues.¹³

Recall the motivations for (1a) rehearsed early in the paper:

- (i) We have a reasonable preference for sound arguments over merely valid arguments.
- (ii) A 'good reason' for accepting a conclusion must be a truth.
- (iii) Luck is apparently involved, even when reasoning validly, in reaching a true conclusion from a false premise (compared with the *guarantee* of a true conclusion when reasoning validly from a truth). The luck involved is plausibly thought to be destructive of knowledge.
- (iv) Gettier cases show that though reasoning through a falsehood isn't a necessary part of a Gettier case, reasoning through a relevant falsehood is sufficient for Gettier-eliminating the conclusion.

Points (i) and (ii) on their own are not substantive points. They seem to state a preference for (1a) but they provide no serious motivation for it. What force might be contained in these points is likely found in points (iii) and (iv) which are clearly more serious attempts to motivate (1a).¹⁴

The motivation coming from points (iii) and (iv) may well be the same. One possible diagnosis of what goes wrong in a Gettier case that involves reasoning through a falsehood (eg, the dog/yard/squirrel case) is that in reasoning through the falsehood in the case, one puts oneself in a situation where (knowledge-destroying) *luck* is needed to get back on the path to truth. That common explanation of this sort of Gettier scenario suggests that points (iii) and (iv) are intimately connected.

This point does not downplay the significance of the Gettier issue in raising worries about my rejection of (1a). One accepting my 'knowledge from falsehood' position does owe an alternative explanation of what has gone wrong in Gettier cases involving reasoning through a falsehood. One can't, obviously enough, explain the lack of knowledge simply by saying, 'the agent reasoned through a falsehood' because the 'knowledge from falsehood' position implies that such reasoning is consistent with inferential knowledge. This is indeed an important and complex issue. I have two points to make about the issue now.

First, though one taking my 'knowledge from falsehood' position does owe an explanation of what goes wrong in falsehood-involving Gettier cases, one taking my position does not have a special burden here. After all, the proposed counterexamples to (1a) should, whether ultimately judged successful or not, significantly worry those who take reasoning through a falsehood to be sufficient for generating a Gettier case. Philosophers on both sides of this overall debate need to reexamine the connection between falsehood and Gettier cases. Second, though the Gettier issues are indeed important and interesting, it seems that the relevant Gettier issues fold back into the issues raised by point (iii)

above. If I can, in reply to point (iii), distinguish cases in which inference through a relevant falsehood generates or constitutes knowledge-destroying luck from cases in which it does not, the account will thereby separate out the 'no-knowledge' Gettier falsehood cases from the 'knowledge from falsehood' cases. I see no way to immediately provide a full account of this distinction. But I can, as promised, sketch and motivate my preferred strategy for dealing with this issue.

The key point in my sketch of a strategy made one brief appearance earlier in the paper. Recall that in each of the counterexamples to (1a) the inferential path from the falsehood to the known conclusion did not seem especially shaky or precarious. Indeed, the path from falsehood to truth in each example seemed remarkably stable and secure. It is no accident that the falsehood 'there are 53 people at my talk' gets me to the truth 'my 100 handout copies are sufficient' nor does it seem that I was epistemically lucky in concluding 'the President is not in Brussels' upon coming to (falsely) believe that he is in Utah based on CNN's accurate report of his 'speaking to supporters in Utah'. Given that both Fayetteville airports are far from Providence, Kim's path to knowledge that Hill will miss his party is quite secure given the evidence of Hill's statement that he is 'flying to Fayetteville'.

This notion of 'stability' that I am opposing with 'luck' is still at the level of metaphor and for now that's where it will have to stay. I do think that the main point is quite suggestive: falsehoods of a certain sort (or, better, of a certain sort in certain *overall* contexts) are well suited to take us to truth. Consider what is probably one (but not the only) sort of common situation in which the path from falsity to truth could not be happily described as 'lucky': approximation. One who falsely believes that pi rounded to the 11th decimal place is 3.14159265358 and uses this mistaken belief in reasoning about circumferences can surely be credited with knowing (with relevant further detail filled in) that the circumference of a one meter diameter circle is 'greater than 3.14 meters'.¹⁵

I'm not, I emphasize, claiming that approximation is what's involved in all possible counterexamples to (1a). It's one important feature involved in some of the examples and it's one that helpfully illustrates the mistake involved in thinking that all inferences through falsehoods involve knowledge-destroying luck. There is no hint of problematic 'luck' in the above inference to the circumference of the one meter circle.¹⁶

My sketch of a response to the initial motivations for (1a) and for articulating an account of when knowledge does and does not arise from falsehood is now in view. I emphasize the strength of the clear counterexamples to (1a). I identify, metaphorically for now, an 'anti-luck' component that all the apparent counterexamples share: 'stability' in the path from the falsehood to the known conclusion. And I illustrate one general and common form of reasoning involving falsehoods that one could profitably defend by appeal to this informal notion of stability. Filling in the details of this story is, as I've emphasized, in some ways a complex matter. The necessity of this project and the inevitability

of its success are, however, strongly suggested by the central cases illustrating that we sometimes get knowledge from falsehood.¹⁷

Notes

1. (1) is used in a variety of defenses of 'knowledge foundationalism' and in other settings as well.
2. (1a) has been endorsed by, among others, Aristotle, Russell, Frege and many living philosophers.
3. For simplicity, and because it might be correct, let's assume that the case of two overdetermining inferential paths to a conclusion appearing to generate knowledge despite an involved falsehood in exactly one of the paths can be diagnosed in the same way as in the overdetermining premises case above (about the Rutgers philosophers of mind). In this case, rule that the path with the falsehood is 'irrelevant'.
4. Alternatively: think of the conclusion in question as being epistemically sustained by a continuous chanting of the premise of one's inferential argument combined with mental affirmation of the suppressed conditional.
5. Points two and three may well be only elaborations of part of the content of the first point.
6. Truth of the conclusion is of course a necessary condition for knowledge of the conclusion.
7. Notice that another consequence of (1) above is this claim:
(1b) Inferential knowledge of a conclusion requires warranted relevant premises. As an aside, notice that in light of this first counterexample to (1a), (1b) seems to fall quickly. Some accept that warrant entails truth. If it does, the falsity of (1a) implies the falsity of (1b). Whether or not warrant entails truth, (1b) falls to an example that parallels the clock example just provided. For the parallel example, leave everything in the initial example the same except this: I have indeed slept 3 ½ hours (we're at 2:30am for the first time). My premise belief is true (but unwarranted due to the relevance of '4 ½ hours') and yet once again I clearly know my conclusion. I'll say no more about (1b) except to note that it has played a non-trivial role in many 'pro-foundationalist' arguments and that it clearly cannot serve in that capacity.
8. As my luggage has informed me, there are two 'Fayetteville' airports in the USA. For a counterexample to (1b) add that unbeknownst to Kim and Hill, Hill's flight will divert to Fayetteville, Arkansas on Saturday.
9. One might make this response for some particular cases and combine this with some other response to the cases conceded to contain a known conclusion.
10. I did this in rejecting the relevance of the premise 'McGinn is in the room' in the argument for the conclusion that 'at least one Rutgers philosopher of mind is present'.
11. Again, one might mix and match, combining more than one of these approaches and also possibly combining this sort of response with an application of the earlier 'reject knowledge' response mentioned a moment ago.

12. This assumption is controversial because it is a claim about *all* members of the class of counterexamples. Of course it's clear that in some of the counterexamples the suggested truth would be at least dispositionally believed. I see no reason to believe that this will be true of all cases in this class.
13. I intend to take up these additional issues on another occasion.
14. It is unlikely that any positive motivation could really come from point (ii). As I'll discuss shortly, approximations known to be false can surely provide 'good reasons' for accepting conclusions in a variety of situations.
15. The 12th decimal place is occupied by '9' and so we need to round up to 3.14159265359.
16. Part of the filling in of the case would be an explanation of why the agent involved would (reasonably) never make judgments of the circumference of the circles he's dealing with out beyond, say, the 6th decimal point. These points about approximation clearly apply to counterexamples #3 and #5. Perhaps they apply to #1 as well.
17. I presented some ideas from this paper in talks at Calvin College and at Notre Dame. Both audiences provided insightful feedback. I especially benefitted from comments from Joshua Armstrong, Robert Audi, Marian David, Michael DePaul, Stephen Grimm, Matt Halteman, Anja Jauernig, Jaegwon Kim, William Ramsey, Michael Schweiger, Leopold Stubenberg and Stephen Wykstra. Robert Audi and EJ Coffman have been extremely generous with feedback and assistance during the drafting of the written version of the paper. Discussions with Coffman, who commented at a conference on related work from Claudio de Almeida and Peter Klein, informed me that others were working on this topic and prompted me to write up my notes for publication.